Estimating Principal Components of Covariance Matrices Using the Nyström Method

Nicholas Arcolano and Patrick J. Wolfe

arXiv:1111.6926v2 [stat.AP] 30 Nov 2011

***Abstract*—Covariance matrix estimates are an essential part of many signal processing algorithms, and are often used to determine a low-dimensional principal subspace via their spectral decomposition. However, exact eigenanalysis is computationally intractable for sufficiently high-dimensional matrices, and in the case of small sample sizes, sample eigenvalues and eigenvectors are known to be poor estimators of their population counterparts. To address these issues, we propose a covariance estimator that is computationally efficient while also performing shrinkage on the sample eigenvalues. Our approach is based on the Nyström method, which uses a data-dependent orthogonal projection to obtain a fast low-rank approximation of a large positive semidefinite matrix. We provide a theoretical analysis of the error properties of our estimator as well as empirical results, including examples of its application to adaptive beamforming and image denoising.**

***Index Terms*—low-rank approximation, covariance shrinkage, Nyström extension, adaptive beamforming, image denoising**

I. Introduction

THE need to determine a principal subspace containing some signal of interest arises in many areas of signal processing, including beamforming [1], speech processing [2], and source separation [3]. Typically, subspace estimation involves computing the spectral decomposition of a covariance matrix that has been estimated using a set of observed data vectors. The estimator most commonly used in practice is the *sample covariance*, often preferred because it is simple to compute and has well-understood theoretical properties.

When solving the subspace estimation problem, one faces two critical challenges. The first is that for p-dimensional data, the computational cost of the full spectral decomposition scales as $O(p^3)$. In the case of high-dimensional data sets, or when the subspace estimation problem needs to be solved many times, obtaining the eigenvalues and eigenvectors of the sample covariance becomes a computational bottleneck. For this reason, algorithms have been developed to obtain approximate solutions to the eigenvalue problem [4], [5].

The second challenge is that the eigenvalues and eigenvectors the sample covariance matrix are known to be be poor estimates of the true eigenvalues and eigenvectors, especially when operating in high dimensions with limited observations [6], [7]. In particular, the sample eigenvalues are known to be over-dispersed (relative to the true spectrum), and many researchers have focused on developing *shrinkage estimators* that yield improved estimation results [8]–[11].

This work is sponsored by the United States Air Force under contract FA8721-05-C-0002. Opinions, interpretations, recommendations, and conclusions are those of the authors and are not necessarily endorsed by the United States Government.

Instead of addressing these challenges separately, we propose that by solving them concurrently, one can perform both tasks at a reduced computational cost. To this end, we develop an estimator based on the *Nyström method* [12]–[14], a matrix approximation technique that uses a data-dependent orthogonal projection to approximate a positive semidefinite matrix. This approach leads to an estimator that not only admits computationally efficient spectral analysis, but also shrinks the eigenvalues of the sample covariance.

We begin by formulating the covariance estimation problem and reviewing the Nyström method for matrix approximation. We then develop its use as a covariance estimator, including a study of its error characteristics. We conclude with examples of the use the Nyström covariance estimator in two practical applications: adaptive beamforming and image denoising.

II. The Covariance Estimation Problem

Let $\mathbf{X}$ be a $p \times n$ matrix whose columns $\mathbf{x}_1, \ldots, \mathbf{x}_n$ are independent and identically distributed (i.i.d.) samples from an unknown p-variate distribution. Throughout the following, we assume $\mathbf{x}_1, \ldots, \mathbf{x}_n$ have zero mean and a finite covariance, denoted by the $p \times p$ positive semidefinite matrix $\mathbf{\Sigma}$.[1]

The basic problem of covariance estimation is straightforward: given $\mathbf{X}$, we wish to construct an estimator $\widehat{\mathbf{\Sigma}}$ of $\mathbf{\Sigma}$. As a function of random data, $\widehat{\mathbf{\Sigma}}$ is itself random, and thus its performance as an estimator is best understood through its statistical properties (conditional on the true covariance). In particular, we will be concerned with the *bias matrix*

$$\mathbf{B}\big(\widehat{\mathbf{\Sigma}} \,\big|\, \mathbf{\Sigma}\big) \equiv \mathbf{\Sigma} - \mathbb{E}\big(\widehat{\mathbf{\Sigma}} \,\big|\, \mathbf{\Sigma}\big),$$

and the *mean squared error* (MSE)

$$\mathrm{MSE}\big(\widehat{\mathbf{\Sigma}} \,\big|\, \mathbf{\Sigma}\big) \equiv \mathbb{E}\left(\left\|\mathbf{\Sigma} - \widehat{\mathbf{\Sigma}}\right\|^2 \,\middle|\, \mathbf{\Sigma}\right),$$

where $\|\cdot\|$ is a suitable matrix norm. A common choice of norm is the Frobenius norm, defined for a real matrix $\mathbf{A}$ as $\|\mathbf{A}\|_F = [\mathrm{tr}(\mathbf{A}^T\mathbf{A})]^{1/2}$. This norm is used throughout the covariance estimation literature [11], [15], and will be the primary one featured in our analysis.

The Sample Covariance

The most common covariance estimator is the *sample covariance matrix*,

$$\mathbf{S} = \frac{1}{n}\sum_{i=1}^{n} \mathbf{x}_i\mathbf{x}_i^T = \frac{1}{n}\mathbf{X}\mathbf{X}^T.$$

[1] We refer to a $p \times p$ matrix $\mathbf{\Sigma}$ as *positive semidefinite* (denoted $\mathbf{\Sigma} \succeq 0$) if it is symmetric and $\mathbf{x}^T\mathbf{\Sigma}\mathbf{x} \geq 0$ for all $\mathbf{x} \in \mathbb{R}^n$, and as *positive definite* (denoted $\mathbf{\Sigma} \succ 0$) if it is positive semidefinite with $\mathbf{x}^T\mathbf{\Sigma}\mathbf{x} = 0$ if and only if $\mathbf{x} = \mathbf{0}$.

This estimator has a number of qualities that make it a popular choice among practitioners. For example, it is unbiased, and its computational cost of $O(p^2 n)$ is not excessively expensive. When $\mathbf{x}_1, \ldots, \mathbf{x}_n$ are i.i.d. samples from the p-variate normal distribution with mean $\mathbf{0}$ and positive definite covariance $\mathbf{\Sigma}$—denoted $\mathcal{N}_p(\mathbf{0}, \mathbf{\Sigma})$—it corresponds to the maximum-likelihood estimator of $\mathbf{\Sigma}$ given the data. We can also compute its MSE with respect to the Frobenius norm, given by

$$\mathrm{MSE}\left(\mathbf{S} \mid \mathbf{\Sigma}\right) \equiv \frac{1}{n}\left[\mathrm{tr}\left(\mathbf{\Sigma}^2\right) + \mathrm{tr}^2\left(\mathbf{\Sigma}\right)\right]. \tag{1}$$

Despite these convenient properties, there are characteristics of $\mathbf{S}$ that make it unsuitable for many applications. For example, if $n < p$, then $\mathbf{S}$ is guaranteed to be rank-deficient; even if we have prior reason to believe that $\mathbf{\Sigma}$ is invertible, its estimate will be singular. Another issue is the over-dispersion of the sample eigenvalues. Let $\lambda_1(\mathbf{S}), \ldots, \lambda_p(\mathbf{S})$ denote the eigenvalues of $\mathbf{S}$ in nonincreasing order. For fixed p and large n, the sample eigenvalues are reasonable estimators of the true eigenvalues of $\mathbf{\Sigma}$, and it can be shown [16] that as $n \to \infty$ with p fixed, $\lambda_i(\mathbf{S})$ converges almost surely to $\lambda_i(\mathbf{\Sigma})$ for $i = 1, \ldots, p$. However, when p is allowed to grow with n (keeping the ratio n/p fixed), results such as the celebrated *Marčkenko-Pastur law* suggest that the sample eigenvalues are not effective estimators, and fail to converge to the true eigenvalues [6], [7].

Shrinkage Covariance Estimators

In the context of covariance estimation, *shrinkage estimation* involves compensating for the known over-dispersion of sample eigenvalues, in order to improve error performance and numerical stability. One common approach to covariance shrinkage is to preserve the sample eigenvectors, but alter the sample eigenvalues to improve error performance with respect to a given loss function [8], [10]. Although shrinkage estimators of this form often come with analytical guarantees regarding error performance, they do so at the cost of operating on the spectral decomposition of $\mathbf{S}$ directly. Consequently, such estimators are impractical for large data sets.

Another approach is to construct the estimate $\widehat{\mathbf{\Sigma}}$ as a linear combination of $\mathbf{S}$ and a known positive definite matrix [9], [11]. For example, the Ledoit-Wolf estimator of [11] takes the form

$$\widehat{\mathbf{\Sigma}} = \alpha\mathbf{S} + \beta\mathbf{I}_p, \tag{2}$$

where the optimal shrinkage coefficients $\alpha, \beta \geq 0$ are given by

$$(\alpha^*, \beta^*) = \underset{(\alpha,\beta)}{\arg\min} \left\|(\alpha\mathbf{S} + \beta\mathbf{I}_p) - \mathbf{\Sigma}\right\|_F^2.$$

Since the optimal coefficients can only be obtained analytically in the case where $\mathbf{\Sigma}$ is known, the authors instead develop consistent estimators for α and β based on the data. Note that while estimators such as (2) have lower computational demands than those that operate on the eigenvalues of $\mathbf{S}$ directly, once we have obtained an estimate $\widehat{\mathbf{\Sigma}}$, we still must pay the full cost of $O(p^3)$ if we wish to obtain its principal components.

Covariance Estimation Using the Nyström Method

Let $I \subseteq \{1, \ldots, p\}$ be a set of k indices. The Nyström covariance estimator [14] takes the form

$$\widehat{\mathbf{\Sigma}}\,(I) = \frac{1}{n}\mathbf{X}\,\mathbf{P}(I)\,\mathbf{X}^T, \tag{3}$$

where $\mathbf{P}\,(I)$ represents an orthogonal projection onto the subspace of $\mathbb{R}^n$ spanned by the k rows of $\mathbf{X}$ specified by the indices in I.

Defining $\widehat{\mathbf{\Sigma}}$ in this fashion serves two important purposes. First, we will show that (3) is equivalent to the *Nyström approximation* of the sample covariance, an established method for low-rank approximation of positive semidefinite matrices [12], [13]. A primary advantage of this method is its computational efficiency, as one can obtain the k principal eigenvalues and eigenvectors of $\widehat{\mathbf{\Sigma}}$ for a cost that scales linearly in p and n. Second, the projection $\mathbf{P}$ shrinks the singular values of the data, serving to counteract over-dispersion of eigenvalues in the sample covariance.

III. The Nyström Method

The Nyström method is a classical technique for obtaining numerical solutions to eigenfunction problems. When applied to matrices, it can be used to construct a low-rank approximation of a positive semidefinite matrix as follows.

Let $\mathbf{Q}$ be a $p \times p$ positive semidefinite matrix, represented in block form as

$$\mathbf{Q} = \begin{bmatrix} \mathbf{Q}_{11} & \mathbf{Q}_{12} \\ \mathbf{Q}_{12}^T & \mathbf{Q}_{22} \end{bmatrix},$$

where $\mathbf{Q}_{11}$ is $k \times k$. The Nyström approximation of $\mathbf{Q}$ preserves $\mathbf{Q}_{11}$ and $\mathbf{Q}_{12}$ while approximating $\mathbf{Q}_{22}$ by its *Nyström extension*:

$$\widehat{\mathbf{Q}} \equiv \begin{bmatrix} \mathbf{Q}_{11} & \mathbf{Q}_{12} \\ \mathbf{Q}_{12}^T & \mathbf{Q}_{12}^T\mathbf{Q}_{11}^+\mathbf{Q}_{12} \end{bmatrix}, \tag{4}$$

where $\mathbf{Q}_{11}^+$ denotes the Moore-Penrose pseudoinverse of $\mathbf{Q}_{11}$. Since the approximation reconstructs $\mathbf{Q}_{11}$ and $\mathbf{Q}_{12}$ exactly, the approximation error $\mathbf{Q} - \widehat{\mathbf{Q}}$ is characterized entirely by the *Schur complement* of $\mathbf{Q}_{11}$ in $\mathbf{Q}$,

$$\overline{\mathbf{Q}}_{11} \equiv \mathbf{Q}_{22} - \mathbf{Q}_{12}^T\mathbf{Q}_{11}^+\mathbf{Q}_{12}.$$

If we view $\mathbf{Q}$ as the outer product of an underlying data matrix, an alternative way to characterize the Nyström approximation is as a function of an orthogonal projection. Let $\mathbf{X}$ be a $p \times n$ matrix, partitioned as

$$\mathbf{X} = \begin{bmatrix} \mathbf{Y} \\ \mathbf{Z} \end{bmatrix} \tag{5}$$

where $\mathbf{Y}$ is $k \times n$ and $\mathbf{Z}$ is $(p-k) \times n$, and let

$$\mathbf{Q} = \mathbf{X}\mathbf{X}^T = \begin{bmatrix} \mathbf{Y}\mathbf{Y}^T & \mathbf{Y}\mathbf{Z}^T \\ \mathbf{Z}\mathbf{Y}^T & \mathbf{Z}\mathbf{Z}^T \end{bmatrix}.$$

We then define the $n \times n$ symmetric idempotent matrix

$$\mathbf{P} \equiv \mathbf{Y}^T\left(\mathbf{Y}^T\right)^+ = \mathbf{Y}^T\left(\mathbf{Y}\mathbf{Y}^T\right)^+\mathbf{Y}, \tag{6}$$

which represents an orthogonal projection onto the subspace of $\mathbb{R}^n$ spanned by the k rows of $\mathbf{Y}$. We obtain the same

expression as in (4) by approximating $\mathbf{X}$ with its projection $\mathbf{XP}$:

$$\widehat{\mathbf{Q}} = \mathbf{XP}(\mathbf{XP})^T = \mathbf{XPX}^T = \begin{bmatrix} \mathbf{YY}^T & \mathbf{YZ}^T \\ \mathbf{ZY}^T & \mathbf{ZY}^T(\mathbf{YY}^T)^+\mathbf{YZ}^T \end{bmatrix}.$$

This interpretation illustrates the low-rank nature of $\widehat{\mathbf{Q}}$, as we must have $\operatorname{rank}(\widehat{\mathbf{Q}}) \leq \operatorname{rank}(\mathbf{P}) \leq k$. It also highlights the fact that the approximation need not be restricted to the first k rows of $\mathbf{X}$; we may instead choose to construct $\mathbf{P}$ based on *any* subset of k rows. Since different subsets typically yield different approximations, we can view $\widehat{\mathbf{Q}}$ as a function of a set of k indices $I \subseteq \{1, \ldots, p\}$. (Throughout the article, we we will use $\mathbf{A}_{IJ}$ to denote the submatrix of a matrix $\mathbf{A}$ whose rows and columns are specified by respective index sets I and J, and define $\mathbf{A}_I \equiv \mathbf{A}_{II}$.)

The problem of selecting a suitable subset for Nyström approximation is one that has received significant attention in the literature. Although a number of efforts have focused on developing advanced subset selection methods [17], [18], for many applications satisfactory performance can be achieved simply by choosing I randomly with uniform probability [12], [13]. This simpler approach has the added benefit of enhancing the computational gains associated with the Nyström method, and will be the strategy employed for the experiments in Sections V and VI.

IV. The Nyström Covariance Estimator

We proceed by formally defining the Nyström covariance estimator, after which we derive expressions for its bias and MSE. We then discuss its eigenvalue shrinkage properties and derive expressions for its eigenvalues and eigenvectors.

Definition 1 (Nyström covariance estimator)**.** *Let $\mathbf{X}$ be a $p \times n$ matrix whose columns $\mathbf{x}_1, \ldots, \mathbf{x}_n$ are i.i.d. random vectors such that $\mathbb{E}(\mathbf{x}_i) = \mathbf{0}$ and $\mathbb{E}(\mathbf{x}_i\mathbf{x}_i^T) = \mathbf{\Sigma}$ for $i = 1, \ldots, n$. Let $\mathbf{r}_1, \ldots, \mathbf{r}_p$ denote the rows of $\mathbf{X}$. Given a k-subset $I \subseteq \{1, \ldots, p\}$, we define the Nyström covariance estimator of $\mathbf{\Sigma}$ as*

$$\widehat{\mathbf{\Sigma}}(I) \equiv \frac{1}{n}\mathbf{X}\,\mathbf{P}(I)\,\mathbf{X}^T,$$

where $\mathbf{P}(I)$ represents an orthogonal projection onto the subspace of $\mathbb{R}^n$ spanned by the set of vectors $\{\mathbf{r}_i : i \in I\}$.

As previously discussed, $\widehat{\mathbf{\Sigma}}(I)$ is a function of an index set $I \subseteq \{1, \ldots, p\}$, and thus error performance will conditional on I. Although viewed here as an estimator of $\mathbf{\Sigma}$, the Nyström covariance estimator could be interpreted as the Nyström approximation of the sample covariance $\mathbf{S}$. When $\operatorname{rank}(\mathbf{P}(I)) = \operatorname{rank}(\mathbf{X}) \leq \min(p, n)$, this approximation is exact, and we have $\widehat{\mathbf{\Sigma}}(I) = \mathbf{S}$.

Error Statistics

Assume now that the columns of $\mathbf{X}$ are drawn independently from a p-variate normal distribution with zero mean and covariance $\mathbf{\Sigma} \succ 0$. In this case, we can derive analytical expressions for the bias and expected square error of the Nyström covariance estimator. We begin by computing the expected value of $\widehat{\mathbf{\Sigma}}(I)$, after which the bias matrix follows as a corollary.

Theorem 1 (Expected value of Nyström covariance estimator)**.** *Let $\mathbf{X}$ be a $p \times n$ matrix whose columns $\mathbf{x}_1, \ldots, \mathbf{x}_n$ are i.i.d. random samples from $\mathcal{N}_p(\mathbf{0}, \mathbf{\Sigma})$. Let $\widehat{\mathbf{\Sigma}}(I)$ be the Nyström covariance estimator of $\mathbf{\Sigma}$ given a k-subset $I \subseteq \{1, \ldots, p\}$, and define $J = \{1, \ldots, p\} \setminus I$. Then, $\big[\mathbb{E}\big(\widehat{\mathbf{\Sigma}}(I)\big)\big]_I = [\mathbf{\Sigma}]_I$, $\big[\mathbb{E}\big(\widehat{\mathbf{\Sigma}}(I)\big)\big]_{IJ} = \big[\mathbb{E}\big(\widehat{\mathbf{\Sigma}}(I)\big)\big]_{JI}^T = [\mathbf{\Sigma}]_{IJ}$, and*

$$\Big[\mathbb{E}\big(\widehat{\mathbf{\Sigma}}(I)\big)\Big]_J = \frac{k}{n}\mathbf{\Sigma}_J + \frac{n-k}{n}\mathbf{\Sigma}_{IJ}^T\mathbf{\Sigma}_J^{-1}\mathbf{\Sigma}_{IJ}\,.$$

Proof: Without loss of generality, let $I = \{1, \ldots, k\}$ and $J = \{k+1, \ldots, p\}$. Partitioning $\mathbf{X}$ as in (5), the Nyström covariance estimate is given by

$$\widehat{\mathbf{\Sigma}} = \frac{1}{n}\mathbf{XPX}^T = \frac{1}{n}\begin{bmatrix} \mathbf{YY}^T & \mathbf{YZ}^T \\ \mathbf{ZY}^T & \mathbf{ZPZ}^T \end{bmatrix},$$

where $\mathbf{P}$ represents an orthogonal projection onto the span of the rows of $\mathbf{Y}$. By construction, we have

$$\mathbb{E}\big(\tfrac{1}{n}\mathbf{YY}^T\big) = \mathbf{\Sigma}_I\,,$$

and

$$\mathbb{E}\big(\tfrac{1}{n}\mathbf{YZ}^T\big) = \big[\mathbb{E}\big(\tfrac{1}{n}\mathbf{ZY}^T\big)\big]^T = \mathbf{\Sigma}_{IJ}\,,$$

and thus we need only compute $\mathbb{E}\big(\frac{1}{n}\mathbf{ZPZ}^T\big)$. To perform this calculation, consider the nested expectation

$$\mathbb{E}\big(\mathbf{ZPZ}^T\big) = \mathbb{E}_Y\Big[\mathbb{E}\big(\mathbf{ZPZ}^T \,|\, \mathbf{Y}\big)\Big].$$

Using standard properties of conditional distributions of normal random vectors, one can show that given $\mathbf{Y}$, the columns $\mathbf{z}_1, \ldots, \mathbf{z}_n$ of $\mathbf{Z}$ are independent and normally distributed as

$$\mathbf{z}_i \,|\, \mathbf{Y} \sim \mathcal{N}_{p-k}\left(\boldsymbol{\mu}_{Z|Y}^{(i)}, \mathbf{\Sigma}_{Z|Y}\right),$$

where

$$\boldsymbol{\mu}_{Z|Y}^{(i)} = \mathbf{\Sigma}_{IJ}^T\mathbf{\Sigma}_I^{-1}\mathbf{y}_i\,,$$

and

$$\mathbf{\Sigma}_{Z|Y} = \mathbf{\Sigma}_J - \mathbf{\Sigma}_{IJ}^T\mathbf{\Sigma}_I^{-1}\mathbf{\Sigma}_{IJ} = \overline{\mathbf{\Sigma}}_I.$$

Given these distributions, evaluating $\mathbb{E}\big(\mathbf{ZPZ}^T \,|\, \mathbf{Y}\big)$ is a matter of performing standard moment calculations using the properties of normal random variables. For convenience, we apply a result from [19, Theorem 2.3.5] for normal random matrices, which states that for a random $p \times n$ matrix $\mathbf{X}$ whose columns $\mathbf{x}_1, \ldots, \mathbf{x}_n$ are distributed as $\mathbf{x}_i \sim \mathcal{N}_p(\boldsymbol{\mu}_i, \mathbf{\Sigma})$, if $\mathbf{A}$ is a constant $p \times p$ matrix, then

$$\mathbb{E}\big(\mathbf{XAX}^T\big) = \operatorname{tr}(\mathbf{A})\,\mathbf{\Sigma} + \mathbf{MAM}^T,$$

where $\mathbf{M} = [\boldsymbol{\mu}_1 \;\cdots\; \boldsymbol{\mu}_n]$. Thus,

$$\mathbb{E}\big(\mathbf{ZPZ}^T \,|\, \mathbf{Y}\big) = k\overline{\mathbf{\Sigma}}_I + \mathbf{\Sigma}_{IJ}^T\mathbf{\Sigma}_I^{-1}\mathbf{YPY}^T\mathbf{\Sigma}_I^{-1}\mathbf{\Sigma}_{IJ}\,,$$

and

$$\begin{aligned}\mathbb{E}\big(\mathbf{ZPZ}^T\big) &= \mathbb{E}_Y\Big[\mathbb{E}\big(\mathbf{ZPZ}^T \,|\, \mathbf{Y}\big)\Big] \\ &= \mathbb{E}_Y\Big[k\overline{\mathbf{\Sigma}}_I + \mathbf{\Sigma}_{IJ}^T\mathbf{\Sigma}_I^{-1}\mathbf{YPY}^T\mathbf{\Sigma}_I^{-1}\mathbf{\Sigma}_{IJ}\Big] \\ &= k\mathbf{\Sigma}_J + (n-k)\,\mathbf{\Sigma}_{IJ}^T\mathbf{\Sigma}_I^{-1}\mathbf{\Sigma}_{IJ},\end{aligned}$$

where the final equality follows from $\mathbb{E}_Y\big(\mathbf{Y}\mathbf{P}\mathbf{Y}^T\big) = n\boldsymbol{\Sigma}_I$. Dividing by n yields the desired result. □

Corollary 1 (Bias of Nyström covariance estimator [14])**.** *Let $\mathbf{X}$ be a $p \times n$ matrix whose columns $\mathbf{x}_1, \dots, \mathbf{x}_n$ are i.i.d. random samples from $\mathcal{N}_p(\mathbf{0}, \boldsymbol{\Sigma})$. Let $\widehat{\boldsymbol{\Sigma}}(I)$ be the Nyström covariance estimator of $\boldsymbol{\Sigma}$ given a k-subset $I \subseteq \{1, \dots, p\}$, and define $J = \{1, \dots, p\} \setminus I$. Then the bias matrix*

$$\mathbf{B}\big(\widehat{\boldsymbol{\Sigma}}(I) \,\big|\, \boldsymbol{\Sigma}\big) = \boldsymbol{\Sigma} - \mathbb{E}\big(\widehat{\boldsymbol{\Sigma}}(I)\big)$$

satisfies $[\mathbf{B}]_{ij} = 0$ for all $(i, j) \notin J \times J$, and

$$\mathbf{B}_J = \frac{n-k}{n}\overline{\boldsymbol{\Sigma}}_I = \frac{n-k}{n}\Big[\boldsymbol{\Sigma}_J - \boldsymbol{\Sigma}_{IJ}^T\boldsymbol{\Sigma}_J^{-1}\boldsymbol{\Sigma}_{IJ}\Big].$$

Thus, $\widehat{\boldsymbol{\Sigma}}(I)$ is a biased estimator of $\boldsymbol{\Sigma}$, except in the case where the Schur complement $\overline{\boldsymbol{\Sigma}}_I = \mathbf{0}$. Recalling from Section III that this Schur complement also expresses the error between $\boldsymbol{\Sigma}$ and its Nyström approximation, we see that $\widehat{\boldsymbol{\Sigma}}(I)$ cannot be unbiased unless it is equal to the sample covariance.

Theorem 2 (MSE of Nyström covariance estimator [14])**.** *Let $\mathbf{X}$ be a $p \times n$ matrix whose columns $\mathbf{x}_1, \dots, \mathbf{x}_n$ are i.i.d. random samples from $\mathcal{N}_p(\mathbf{0}, \boldsymbol{\Sigma})$. Let $\widehat{\boldsymbol{\Sigma}}(I)$ be the Nyström covariance estimator of $\boldsymbol{\Sigma}$ given a k-subset $I \subseteq \{1, \dots, p\}$, and define $J = \{1, \dots, p\} \setminus I$. Then the mean square error of the Nyström covariance estimator in Frobenius norm is*

$$\begin{aligned}
&\mathbb{E}\left\|\boldsymbol{\Sigma} - \widehat{\boldsymbol{\Sigma}}(I)\right\|_F^2 \\
&\quad= \mathrm{MSE}\left(\mathbf{S} \mid \boldsymbol{\Sigma}\right) + \frac{n-k}{n^2}\Big[(n-k-1)\mathrm{tr}\big(\overline{\boldsymbol{\Sigma}}_I^2\big) - \mathrm{tr}^2\big(\overline{\boldsymbol{\Sigma}}_I\big)\Big] \\
&\quad= \mathrm{MSE}\left(\mathbf{S} \mid \boldsymbol{\Sigma}\right) + \frac{(n-k)^2}{n^2}\Big[\left\|\overline{\boldsymbol{\Sigma}}_I\right\|_F^2 - \mathrm{MSE}\left(\overline{\mathbf{S}}_I \mid \boldsymbol{\Sigma}\right)\Big],
\end{aligned}$$

where $\mathrm{MSE}\left(\mathbf{S} \mid \boldsymbol{\Sigma}\right)$ is the mean square error of the sample covariance estimator given in (1)*, and $\mathrm{MSE}\left(\overline{\mathbf{S}}_I \mid \boldsymbol{\Sigma}\right)$ is the mean square error of the sample covariance estimator of the Schur complement of $\boldsymbol{\Sigma}_I$ in $\boldsymbol{\Sigma}$, given by*

$$\mathrm{MSE}\left(\overline{\mathbf{S}}_I \mid \boldsymbol{\Sigma}\right) \equiv \frac{1}{n-k}\Big[\mathrm{tr}\big(\overline{\boldsymbol{\Sigma}}_I^2\big) + \mathrm{tr}^2\big(\overline{\boldsymbol{\Sigma}}_I\big)\Big].$$

Proof: Let $I = \{1, \dots, k\}$ and $J = \{k+1, \dots, p\}$ without loss of generality. The MSE in Frobenius norm is

$$\begin{aligned}
\mathbb{E}\left\|\boldsymbol{\Sigma} - \widehat{\boldsymbol{\Sigma}}(I)\right\|_F^2 = \mathrm{tr}\big(\boldsymbol{\Sigma}^2\big) &- 2\,\mathrm{tr}\Big[\boldsymbol{\Sigma}\,\mathbb{E}\left(\tfrac{1}{n}\mathbf{X}\mathbf{P}\mathbf{X}^T\right)\Big] \\
&+ \mathrm{tr}\Big[\mathbb{E}\left(\tfrac{1}{n^2}\mathbf{X}\mathbf{P}\mathbf{X}^T\mathbf{X}\mathbf{P}\mathbf{X}^T\right)\Big]. \quad (7)
\end{aligned}$$

Substituting the expression for $\mathbb{E}\left(\frac{1}{n}\mathbf{X}\mathbf{P}\mathbf{X}^T\right)$ from Theorem 1, we have

$$\begin{aligned}
\mathrm{tr}\left[\boldsymbol{\Sigma}\,\mathbb{E}\left(\frac{1}{n}\mathbf{X}\mathbf{P}\mathbf{X}^T\right)\right] &= \mathrm{tr}\big(\boldsymbol{\Sigma}_I^2\big) + 2\,\mathrm{tr}\Big(\boldsymbol{\Sigma}_{IJ}\boldsymbol{\Sigma}_{IJ}^T\Big) + \tfrac{k}{n}\mathrm{tr}\big(\boldsymbol{\Sigma}_J^2\big) \\
&\quad + \tfrac{(n-k)}{n}\mathrm{tr}\Big(\boldsymbol{\Sigma}_J\boldsymbol{\Sigma}_{IJ}^T\boldsymbol{\Sigma}_I^{-1}\boldsymbol{\Sigma}_{IJ}\Big). \quad (8)
\end{aligned}$$

To compute $\mathbb{E}\left(\frac{1}{n^2}\mathbf{X}\mathbf{P}\mathbf{X}^T\mathbf{X}\mathbf{P}\mathbf{X}^T\right)$, let $\mathbf{X}$ be partitioned as in (5), so that

$$\begin{aligned}
\mathrm{tr}\big(\mathbf{X}\mathbf{P}\mathbf{X}^T\mathbf{X}\mathbf{P}\mathbf{X}^T\big) &= \mathrm{tr}\left(\mathbf{Y}\mathbf{Y}^T\mathbf{Y}\mathbf{Y}^T\right) + 2\,\mathrm{tr}\big(\mathbf{Y}\mathbf{Z}^T\mathbf{Z}\mathbf{Y}^T\big) \\
&\quad + \mathrm{tr}\big(\mathbf{Z}\mathbf{P}\mathbf{Z}^T\mathbf{Z}\mathbf{P}\mathbf{Z}^T\big). \quad (9)
\end{aligned}$$

As in the proof of Theorem 1, the expectation of each term can be evaluated using standard properties of normal random vectors. However, we may simplify the analysis using a result from [19, Theorem 2.3.8], which states that for a random $p \times n$ matrix $\mathbf{X}$ whose columns $\mathbf{x}_1, \dots, \mathbf{x}_n$ are distributed as $\mathbf{x}_i \sim \mathcal{N}_p\left(\boldsymbol{\mu}_i, \boldsymbol{\Sigma}\right)$, if $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{C}$ are independent $n \times n$, $p \times p$, and $n \times n$ matrices (respectively), then

$$\begin{aligned}
&\mathbb{E}\left(\mathbf{X}\mathbf{A}\mathbf{X}^T\mathbf{B}\mathbf{X}\mathbf{C}\mathbf{X}^T\right) \\
&\quad= \mathrm{tr}\big(\mathbf{C}^T\mathbf{A}^T\big)\mathrm{tr}\left(\mathbf{B}\boldsymbol{\Sigma}\right)\boldsymbol{\Sigma} + \mathrm{tr}\left(\mathbf{A}\right)\mathrm{tr}\left(\mathbf{C}\right)\boldsymbol{\Sigma}\mathbf{B}\boldsymbol{\Sigma} \\
&\qquad+ \mathrm{tr}\big(\mathbf{A}\mathbf{C}^T\big)\boldsymbol{\Sigma}\mathbf{B}^T\boldsymbol{\Sigma} + \mathrm{tr}\left(\mathbf{C}\right)\mathbf{M}\mathbf{A}\mathbf{M}^T\mathbf{B}\boldsymbol{\Sigma} \\
&\qquad+ \mathbf{M}\mathbf{A}\mathbf{C}^T\mathbf{M}^T\mathbf{B}^T\boldsymbol{\Sigma} + \mathrm{tr}\big(\mathbf{A}\mathbf{M}^T\mathbf{B}\mathbf{M}\mathbf{C}\big)\boldsymbol{\Sigma} \\
&\qquad+ \mathrm{tr}\left(\mathbf{B}\boldsymbol{\Sigma}\right)\mathbf{M}\mathbf{A}\mathbf{C}\mathbf{M}^T + \boldsymbol{\Sigma}\mathbf{B}^T\mathbf{M}\mathbf{A}^T\mathbf{C}\mathbf{M}^T \\
&\qquad+ \mathrm{tr}\left(\mathbf{A}\right)\boldsymbol{\Sigma}\mathbf{B}\mathbf{M}\mathbf{C}\mathbf{M}^T + \mathbf{M}\mathbf{A}\mathbf{M}^T\mathbf{B}\mathbf{M}\mathbf{C}\mathbf{M}^T,
\end{aligned}$$

where $\mathbf{M} = [\boldsymbol{\mu}_1 \;\cdots\; \boldsymbol{\mu}_n]$. We can compute the first term in (9) by applying this formula directly; for the remaining two terms, we must use iterated expectation as we did in the proof of Theorem 1. Evaluating these expectations yields

$$\mathbb{E}\,\mathrm{tr}\left(\mathbf{Y}\mathbf{Y}^T\mathbf{Y}\mathbf{Y}^T\right) = \big(n^2+n\big)\,\mathrm{tr}\big(\boldsymbol{\Sigma}_I^2\big) + n\,\mathrm{tr}^2\big(\boldsymbol{\Sigma}_I\big), \quad (10)$$

$$\begin{aligned}
\mathbb{E}\,\mathrm{tr}\left(\mathbf{Y}\mathbf{Z}^T\mathbf{Z}\mathbf{Y}^T\right) &= \big(n^2+n\big)\,\mathrm{tr}\big(\boldsymbol{\Sigma}_{IJ}\boldsymbol{\Sigma}_{IJ}^T\big) \\
&\quad + n\,\mathrm{tr}\big(\boldsymbol{\Sigma}_I\big)\mathrm{tr}\big(\boldsymbol{\Sigma}_J\big), \quad (11)
\end{aligned}$$

and

$$\begin{aligned}
\mathbb{E}\,\mathrm{tr}\big(\mathbf{Z}\mathbf{P}\mathbf{Z}^T\mathbf{Z}\mathbf{P}\mathbf{Z}^T\big) &= \big(k^2+k\big)\,\mathrm{tr}\big(\overline{\boldsymbol{\Sigma}}_I^2\big) + k\,\mathrm{tr}^2\big(\overline{\boldsymbol{\Sigma}}_I\big) \\
&\quad + 2n\big(k+1\big)\,\mathrm{tr}\big(\overline{\boldsymbol{\Sigma}}_I\mathbf{R}\big) \\
&\quad + 2n\big(k+1\big)\,\mathrm{tr}\big(\overline{\boldsymbol{\Sigma}}_I\big)\mathrm{tr}\big(\mathbf{R}\big) \\
&\quad + \big(n^2+n\big)\,\mathrm{tr}\big(\mathbf{R}^2\big) + k\,\mathrm{tr}^2\big(\mathbf{R}\big), \quad (12)
\end{aligned}$$

where $\mathbf{R} \equiv \boldsymbol{\Sigma}_{IJ}^T\boldsymbol{\Sigma}_I^{-1}\boldsymbol{\Sigma}_{IJ}$. Substituting (8) and (10)–(12) into (7) and simplifying terms yields the result. □

Discussion of Error Results

Given an arbitrary covariance $\boldsymbol{\Sigma}$, we can derive a lower bound for the MSE of the Nyström covariance estimator as follows. First, note that when applied to the spectrum of a positive semidefinite matrix $\mathbf{A}$, the Cauchy-Schwarz inequality implies that $\mathrm{tr}^2\left(\mathbf{A}\right) \leq \mathrm{tr}\big(\mathbf{A}^2\big)\,\mathrm{rank}\left(\mathbf{A}\right)$. For $\mathbf{A} = \overline{\boldsymbol{\Sigma}}_I$, substituting this inequality into the error expression in Theorem 2 yields

$$\begin{aligned}
\mathbb{E}\left\|\boldsymbol{\Sigma} - \widehat{\boldsymbol{\Sigma}}(I)\right\|_F^2 &\geq \mathrm{MSE}\left(\mathbf{S} \mid \boldsymbol{\Sigma}\right) + \tfrac{n-k}{n^2}\left(n-k-1\right)\mathrm{tr}\big(\overline{\boldsymbol{\Sigma}}_I^2\big) \\
&\quad - \tfrac{n-k}{n^2}\mathrm{tr}\big(\overline{\boldsymbol{\Sigma}}_I^2\big)\,\mathrm{rank}\big(\overline{\boldsymbol{\Sigma}}_I^2\big) \\
&= \mathrm{MSE}\left(\mathbf{S} \mid \boldsymbol{\Sigma}\right) + \tfrac{(n-k)(n-p-1)}{n^2}\mathrm{tr}\big(\overline{\boldsymbol{\Sigma}}_I^2\big),
\end{aligned}$$

where $\mathrm{rank}\big(\overline{\boldsymbol{\Sigma}}_I^2\big) = p-k$. Thus, $n \leq p$ is a necessary condition for the MSE of the Nyström covariance estimator to be less than that of the sample covariance.

We can show that equality is achieved for this bound when $\boldsymbol{\Sigma} = \mathbf{I}_p$. In this case, $\overline{\boldsymbol{\Sigma}}_I = \mathbf{I}_{p-k}$ for all k-subsets

$I \subseteq \{1,\dots,p\}$, and thus

$$\begin{aligned}\mathbb{E}\left\|\boldsymbol{\Sigma}-\widehat{\boldsymbol{\Sigma}}(I)\right\|_F^2 &= \mathrm{MSE}\left(\mathbf{S}\mid\boldsymbol{\Sigma}\right)+\tfrac{n-k}{n^2}(n-k-1)(p-k)\\ &\quad-\tfrac{n-k}{n^2}(p-k)^2\\ &= \mathrm{MSE}\left(\mathbf{S}\mid\boldsymbol{\Sigma}\right)+\tfrac{(n-k)(p-k)(n-p-1)}{n^2},\end{aligned}$$

where $\mathrm{MSE}\left(\mathbf{S}\mid\boldsymbol{\Sigma}\right) = (p^2+p)/n$. Since the second term in the above summation is negative if and only if $n-p-1<0$, we see that when estimating the identity covariance, the Nyström estimator achieves better error performance than the sample covariance for any $n \leq p$.

This behavior may seem surprising, especially when noting that $n \leq p$ implies that the sample covariance will be of rank n with probability one, while the Nyström estimator will be of rank $k \leq n$. However, eigenvalue results such as the Marčenko-Pastur theorem indicate that in this regime the sample covariance will be over-dispersed. Even though the Nyström estimator has fewer nonzero eigenvalues than does the sample covariance, the shrinkage it provides on these eigenvalues compensates for its lower rank, resulting in better overall error.

Eigenvalue Shrinkage

The error performance of the Nyström covariance estimator derives from its ability to shrink the over-dispersed eigenvalues of the sample covariance. One way to understand this property is as a consequence of the well-known eigenvalue inequalities of Weyl, a form of which [20, Corollary III.2.3] states that for any $p\times p$ symmetric matrix $\mathbf{A}$ and $p\times p$ positive semidefinite matrix $\mathbf{B}$,

$$\lambda_i\left(\mathbf{A}+\mathbf{B}\right) \geq \lambda_i\left(\mathbf{A}\right), \tag{13}$$

for $i = 1,\dots,p$.

Let $\mathbf{S} \succeq 0$ be a $p\times p$ sample covariance matrix, and let $\widehat{\boldsymbol{\Sigma}}(I)$ be the corresponding Nyström covariance estimator given a k-subset $I \subseteq \{1,\dots,p\}$. As in previous discussions, we may let $I = \{1,\dots,k\}$ without loss of generality. By positive semidefiniteness of the Schur complement, the error matrix

$$\mathbf{E} \equiv \mathbf{S}-\widehat{\boldsymbol{\Sigma}}(I) = \begin{bmatrix}\mathbf{0} & \mathbf{0}\\ \mathbf{0} & \overline{\mathbf{S}}_I\end{bmatrix}$$

is also positive semidefinite. Thus, (13) implies that

$$\lambda_i\big(\widehat{\boldsymbol{\Sigma}}(I)+\mathbf{E}\big) = \lambda_i\big(\mathbf{S}\big) \geq \lambda_i\big(\widehat{\boldsymbol{\Sigma}}(I)\big),$$

for $i = 1,\dots,p$. In other words, given any k-subset, the Nyström covariance estimator shrinks the eigenvalues of the sample covariance toward zero.

Calculation of Eigenvalues and Eigenvectors

We now derive expressions for the eigenvalues and eigenvectors of the Nyström covariance estimator. For notational convenience, given an $m\times n$ matrix $\mathbf{A}$ and a k-subset $I \subseteq \{1,\dots,m\}$, let $\mathbf{A}_{I:}$ denote the $k\times n$ submatrix formed by taking k entire rows of $\mathbf{A}$ as indexed by I. Also, given $r = \mathrm{rank}\left(\mathbf{A}\right)$, we define the *thin singular value decomposition* (thin SVD) of $\mathbf{A}$ as $\mathbf{A} = \mathbf{U}\mathbf{D}\mathbf{V}^T$, where $\mathbf{D}$ is an $r\times r$ diagonal matrix containing the nonzero singular values of $\mathbf{A}$, and $\mathbf{U}$ and $\mathbf{V}$ are $m\times r$ and $n\times r$ matrices whose columns are the corresponding left and right singular vectors.

Using this notation, the eigenvalues and eigenvectors of the Nyström covariance estimator can be expressed as follows.

Theorem 3 (Eigenvalues and eigenvectors of the Nyström covariance estimator)**.** *Let $\mathbf{X}$ be a $p\times n$ matrix whose columns $\mathbf{x}_1,\dots,\mathbf{x}_n$ are i.i.d. random vectors such that $\mathbb{E}(\mathbf{x}_i) = \mathbf{0}$ and $\mathbb{E}(\mathbf{x}_i\mathbf{x}_i^T) = \boldsymbol{\Sigma}$ for $i = 1,\dots,n$. Given a k-subset $I \subseteq \{1,\dots,p\}$, define $J = \{1,\dots,n\}\backslash I$, let $r = \mathrm{rank}\big(\mathbf{X}_{I:}\big)$, and let $\mathbf{X}_{I:}$ have the thin SVD $\mathbf{X}_{I:} = \mathbf{U}_X\mathbf{D}_X\mathbf{V}_X^T$. If $\mathbf{W}$ is the $p\times r$ matrix given by*

$$\begin{aligned}\mathbf{W}_{I:} &= \frac{1}{\sqrt{n}}\,\mathbf{U}_X\mathbf{D}_X,\\ \mathbf{W}_{J:} &= \frac{1}{\sqrt{n}}\,\mathbf{X}_{J:}\mathbf{V}_X,\end{aligned}$$

then $\mathbf{W}\mathbf{W}^T = \widehat{\boldsymbol{\Sigma}}(I)$ is the Nyström covariance estimator of $\boldsymbol{\Sigma}$ given I, and its r nonzero eigenvalues and corresponding eigenvectors are given by $\boldsymbol{\Lambda}^2$ and $\mathbf{U}$, where $\mathbf{W} = \mathbf{U}\boldsymbol{\Lambda}\mathbf{V}^T$ is the thin SVD of $\mathbf{W}$.

Proof: As in previous proofs, we let $I = \{1,\dots,k\}$ and $J = \{k+1,\dots,p\}$ without loss of generality. Let $\mathbf{X}$ be partitioned as in (5), and let $\mathbf{Y} = \mathbf{X}_{I:}$ have the thin spectral decomposition $\mathbf{Y} = \mathbf{U}_Y\mathbf{D}_Y\mathbf{V}_Y^T$. Letting

$$\mathbf{W} = \frac{1}{\sqrt{n}}\begin{bmatrix}\mathbf{U}_Y\mathbf{D}_Y\\ \mathbf{Z}\mathbf{V}_Y\end{bmatrix},$$

we have

$$\begin{aligned}\mathbf{W}\mathbf{W}^T &= \frac{1}{n}\begin{bmatrix}\mathbf{U}_Y\mathbf{D}_Y\\ \mathbf{Z}\mathbf{V}_Y\end{bmatrix}\begin{bmatrix}\mathbf{D}_Y\mathbf{U}_Y^T & \mathbf{V}_Y^T\mathbf{Z}^T\end{bmatrix}\\ &= \frac{1}{n}\begin{bmatrix}\mathbf{U}_Y\mathbf{D}_Y^2\mathbf{U}_Y^T & \mathbf{U}_Y\mathbf{D}_Y\mathbf{V}_Y^T\mathbf{Z}^T\\ \mathbf{Z}\mathbf{V}_Y\mathbf{D}_Y\mathbf{U}_Y^T & \mathbf{Z}\mathbf{V}_Y\mathbf{V}_Y^T\mathbf{Z}^T\end{bmatrix}\\ &= \frac{1}{n}\begin{bmatrix}\mathbf{Y}\mathbf{Y}^T & \mathbf{Y}\mathbf{Z}^T\\ \mathbf{Z}\mathbf{Y}^T & \mathbf{Z}\mathbf{V}_Y\mathbf{V}_Y^T\mathbf{Z}^T\end{bmatrix}.\end{aligned}$$

Noting that the Moore-Penrose pseudoinverse of $\mathbf{Y}^T$ is

$$\big(\mathbf{Y}^T\big)^+ = \mathbf{V}_Y\mathbf{D}_Y^{-1}\mathbf{U}^T,$$

we see that Nyström projection of (6) is equivalent to

$$\mathbf{P} \equiv \mathbf{Y}^T\big(\mathbf{Y}^T\big)^+ = \mathbf{V}_Y\mathbf{D}_Y\mathbf{U}_Y^T\mathbf{U}_Y\mathbf{D}_Y^{-1}\mathbf{V}_Y^T = \mathbf{V}_Y\mathbf{V}_Y^T,$$

and thus $\mathbf{W}\mathbf{W}^T = \widehat{\boldsymbol{\Sigma}}(I)$. Consequently, if $\mathbf{U}\boldsymbol{\Lambda}\mathbf{V}^T$ is the thin SVD of $\mathbf{W}$, then

$$\widehat{\boldsymbol{\Sigma}}(I) = \mathbf{U}\boldsymbol{\Lambda}\mathbf{V}^T\mathbf{V}\boldsymbol{\Lambda}\mathbf{U}^T = \mathbf{U}\boldsymbol{\Lambda}^2\mathbf{U}^T$$

is the (thin) spectral decomposition of $\widehat{\boldsymbol{\Sigma}}(I)$, where $\boldsymbol{\Lambda}^2$ is an $r\times r$ matrix containing the nonzero eigenvalues of $\widehat{\boldsymbol{\Sigma}}(I)$, and $\mathbf{U}$ is a $p\times r$ matrix whose columns are the corresponding eigenvectors. □

For fixed k, the computational complexity of the spectral decomposition of $\widehat{\boldsymbol{\Sigma}}(I)$ is dominated by two operations. The first is the multiplication $\mathbf{X}_{J:}\mathbf{V}_X$, the cost of which scales as $O(p\,n)$; the second is the thin SVD of $\mathbf{W}$, the cost of which scales as $O(p)$. Thus, for n fixed the overall cost of eigenanalysis scales linearly in the dimension p, making Nyström

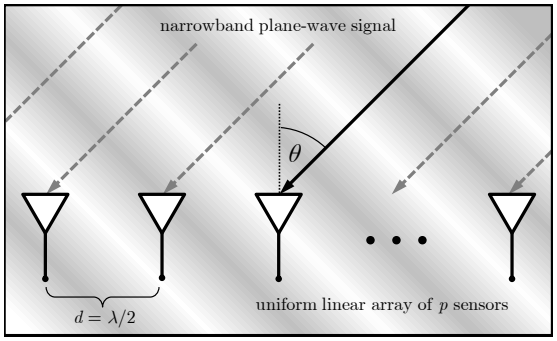

Figure 1. A simple narrowband beamforming model.

covariance estimator an appealing choice in extremely high-dimensional settings.

V. Example Application: Adaptive Beamforming

We conclude our discussion with two examples of the use of Nyström covariance estimation in practical applications. For the first example, we examine the classical signal processing problem of *beamforming* [21]–[24], which involves tailoring the signal response (or "beam pattern") of an array of receiving elements to enhance one's ability to detect a desired signal.

Narrowband Beamforming Model

We adopt a standard beamforming model, illustrated in Figure 1. Consider a collection of p sensing elements, which sample an incoming plane-wave signal at discrete time intervals. We assume the signal of interest is narrowband and that the sensors are placed in a straight line with equal spacing (known as a *uniform linear array*). To avoid aliasing in the spatial sampling of the signal, we assume an element spacing of $d = \lambda_c/2$, where λ_c is the carrier wavelength. In addition, let θ denote the angle between the wave's direction of propagation and a vector normal to the array, referred to as the *angle of arrival*.

Assume that there are $k \leq p$ incoming signals, and let $z_i(t) \in \mathbb{C}$ denote the complex envelope of the i-th signal at the t-th sample time, for $i \in \{1, \dots, k\}$ and $t \in \{1, 2, \dots\}$. Arranging the signal response for all array elements into a vector, the total received signal at the t-th sample time (referred to as a "snapshot") is

$$\mathbf{x}(t) = \sum_{i=1}^{k} \mathbf{a}(\theta_i)\, z_i(t) + \mathbf{n}(t), \tag{14}$$

where $\mathbf{n}(t) \in \mathbb{C}^p$ is additive noise and where $\mathbf{a}(\theta_i) \in \mathbb{C}^p$ represents the amplitude change and phase delay at each sensor as a function of θ_i, the angle of arrival of the i-th signal. Assuming a constant amplitude response across all elements and letting the phase response at the first element be zero, we have

$$[\mathbf{a}(\theta_i)]_l = e^{-j\pi(l-1)\sin(\theta_i)},$$

for $l = 1, \dots, p$ and where $j = \sqrt{-1}$.

Now, let us assume that out of the k incoming signals, only one—say, $z_1(t)$—is of interest to us, and all others are considered interference. In this context, the classical narrowband beamforming problem can be stated as follows: given a collection of n snapshots $\{\mathbf{x}(1), \dots, \mathbf{x}(n)\}$, determine a weight vector $\mathbf{w} \in \mathbb{C}^p$ such that the output of the linear filter (or "beamformer")

$$\hat{z}_1(t) \equiv \mathbf{w}^H \mathbf{x}(t)$$

such that the mean squared error of $z_1(t) - \hat{z}_1(t)$ is small.

For our simple example, let us assume that given the narrowband beamforming model in (14), the signals $z_1(t), \dots, z_m(t)$ and the noise $\mathbf{n}(t)$ are stationary zero-mean Gaussian random processes that are statistically independent across time samples. Let $\sigma_i^2 = \mathbb{E}\big(z_i^2(t)\big)$ denote the expected power of the i-th signal, and let $\sigma_n^2 = \mathbb{E}\big(n^2(t)\big)$ denote the expected noise power. In this case, it can be shown [24], [25] that the optimal beamformer (in the sense of minimum MSE) is given by

$$\mathbf{w}_{\text{opt}} \equiv \min_{\mathbf{w} \in \mathbb{C}^p} \mathbb{E}\left(z_1(t) - \mathbf{w}^H \mathbf{x}(t)\right)^2 = \mathbf{\Sigma}^{-1} \mathbf{a}(\theta_1)\, \sigma_1^2, \tag{15}$$

where

$$\mathbf{\Sigma} = \mathbb{E}\left(\mathbf{x}(t)\, \mathbf{x}^H(t)\right) = \sum_{i=1}^{k} \sigma_i^2\, \mathbf{a}(\theta_i)\, \mathbf{a}^H(\theta_i) + \sigma_n^2\, \mathbf{I}_p\,.$$

Unfortunately, there are a number of barriers to realizing this optimal beamformer in practice. Even if our modeling assumptions hold, we need to know the power σ_1^2 and angle of arrival θ_1 of the signal of interest, as well as the covariances of the interference and noise.

Beamforming Using the Nyström Covariance Estimator

Let assume that σ_1^2 and θ_1 are known. In this case, we might consider approximating $\mathbf{\Sigma}$ by the sample covariance

$$\mathbf{S} = \frac{1}{n} \sum_{t=1}^{n} \mathbf{x}(t)\, \mathbf{x}^H(t),$$

given a collection of n observations. Although this approach will recover the optimal beamformer as $n \to \infty$, in practice the number of samples is bounded, and any assumption of stationarity regarding the signal and interference usually is valid only over a limited time window. This issue is especially problematic when the number of elements p is large, as we may not have enough samples for $\mathbf{S}$ to be well-conditioned (or even invertible). As a result, direct substitution of the sample covariance into (15) is rarely an acceptable solution.

One alternative is to replace $\mathbf{S}$ with a low-rank approximation [23], [26]–[28]. For example, consider the optimal rank-m approximation

$$\mathbf{S}_k^* = \mathbf{U}_k \mathbf{\Lambda}_k \mathbf{U}_k^H,$$

where $\mathbf{\Lambda}_k$ is a $k \times k$ diagonal matrix containing the k largest eigenvalues of $\widehat{\mathbf{\Sigma}}$, and $\mathbf{U}_k$ is a $p \times k$ matrix containing the corresponding eigenvectors. We can define a "projection beamformer" (denoted $\mathbf{w}_{\text{proj}}$) by approximating (15) as

$$\mathbf{w}_{\text{proj}} \equiv \left(\mathbf{S}_k^*\right)^+ \mathbf{a}(\theta_1)\, \sigma_1^2, \tag{16}$$

where

$$\left(\mathbf{S}_k^*\right)^+ = \mathbf{U}_k \mathbf{\Lambda}_k^{-1} \mathbf{U}_k^H.$$

Consequently, the filter output can be expressed as

$$\hat{z}_1(t) = \mathbf{w}_{\text{proj}}^H \mathbf{x}(t) = \sigma_1^2 \mathbf{a}^H(\theta_1) \left(\mathbf{S}_k^*\right)^+ \mathbf{P}\mathbf{x}(t),$$

where $\mathbf{P} \equiv \mathbf{U}_m \mathbf{U}_m^H$ is an orthogonal projection onto the m-dimensional principal subspace of $\mathbf{S}$. Thus, the projection beamformer maps the incoming signal onto a low-rank subspace, and then approximates the behavior of optimal beamformer within this space. If the range of $\mathbf{P}$ is close to the span of $\{\mathbf{a}(\theta_1)z_1(t), \ldots, \mathbf{a}(\theta_m)z_m(t)\}$, any signal power lost in the low-rank projection will be largely due to noise, thus improving overall estimation performance. Of course, the effectiveness of this approach in practice depends on the powers of the signal and interference relative to the noise, as well as the sample size.

The projection beamformer of (16) presents an excellent opportunity to apply the Nyström covariance estimator. Substituting $\widehat{\mathbf{\Sigma}}(I)$ for the optimal low-rank approximation of the sample covariance, we define the "Nyström beamformer" (denoted $\mathbf{w}_{\text{Nyst}}$) as

$$\mathbf{w}_{\text{Nyst}} \equiv \left(\widehat{\mathbf{\Sigma}}(I)\right)^+ \mathbf{a}(\theta_1)\, \sigma_1^2, \tag{17}$$

given a k-subset $I \subseteq \{1, \ldots, p\}$. In contrast to to the projection beamformer, the Nyström beamformer characterizes the signal covariance only over a subset of $k \leq p$ sensors in the array; the rest of the covariance is inferred as a function of observed correlations between elements in this subset and the remaining sensors. The goal of this approach is to achieve performance comparable to that of the projection beamformer, while realizing significant reductions in computational cost.

Experimental Results

To compare the performance of various beamforming approaches, we simulated a uniform linear array with $p = 100$. We considered a case with $k = 7$ signals, where the angle of arrival of desired signal was 10 degrees, and the six interference signals had angles of arrival of -65, -30, -25, 30, 45, and 60 degrees. For all experiments, we assumed a constant noise power of $\sigma_n^2 = 1$ and a constant interference-to-noise ratio (INR) of 20 dB.

Given a signal-to-noise ratio (SNR) for the desired signal, we studied the performance of each method as a function of the number of snapshots n. Our primary measure of performance was the *signal to interference and noise ratio* (SINR) of the estimated signal, defined as

$$\text{SINR} = \frac{\sum_{t=1}^n |\hat{z}_1(t)|^2}{\sum_{t=1}^n |\mathbf{w}^H \mathbf{z}(t)|^2},$$

where $\hat{z}_1(t) = \mathbf{w}^H \mathbf{x}(t)$ and $\mathbf{z}(t)$ is the sum of the received interference and noise at the t-th sample time,

$$\mathbf{z}(t) \equiv \sum_{i=2}^m \mathbf{a}(\theta_i) z_i(t) + \mathbf{n}(t).$$

Figure 2 shows the performance of various beamformers for three different SNR levels (-10 dB, 10 dB, and 30 dB). For each beamformer, we plot the SINR in dB as a function of the number of snapshots, averaged over 1000 experimental trials.

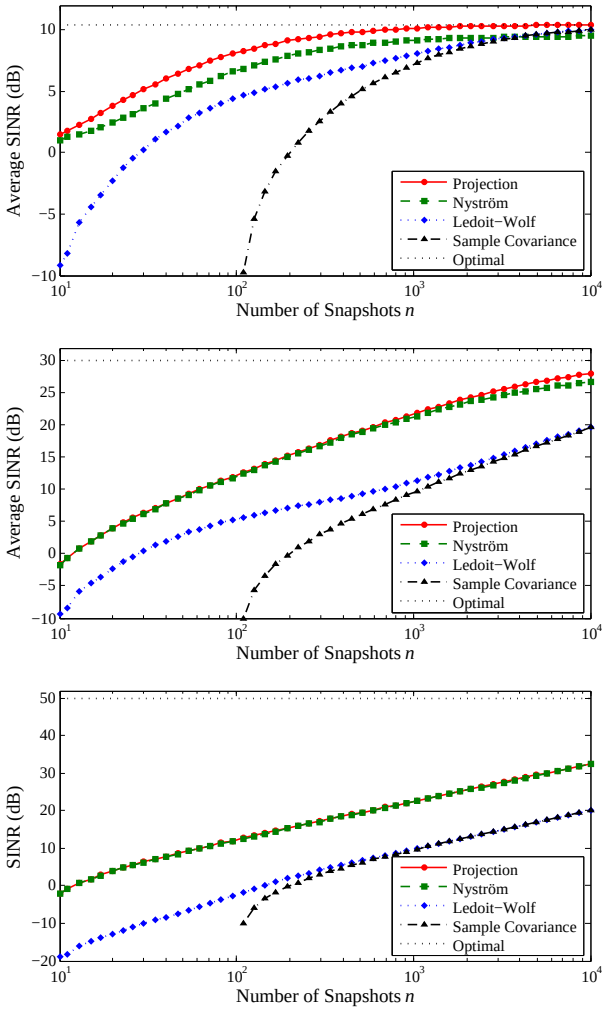

Figure 2. SINR as a function of number of snapshots for various beamforming approaches, given an INR of 20 dB and SNR values of -10 dB (top), 10 dB (middle), and 30 dB (bottom).

Note that the horizontal axis is measured on a logarithmic scale, ranging from $n = 10$ to $n = 10^4$ samples.

Results are shown for projection beamformer of (16) and the Nyström beamformer of (17), where the latter is computed using a uniformly random k-subset $I \subseteq \{1, \ldots, p\}$. For comparison, we include two approximations of the optimal beamformer: one where we have replaced $\mathbf{\Sigma}$ in (15) with the sample covariance, and another where we have substituted the Ledoit-Wolf covariance estimator of [11]. We also show an upper bound given by the theoretical SINR of the optimal beamformer,

$$\text{SINR}_{\text{opt}} = \frac{\mathbb{E}\left|\mathbf{w}_{\text{opt}}^H \mathbf{x}(t)\right|^2}{\mathbb{E}\left|\mathbf{w}_{\text{opt}}^H \mathbf{z}(t)\right|^2} = \frac{\mathbf{w}_{\text{opt}}^H \mathbf{\Sigma} \mathbf{w}_{\text{opt}}}{\mathbf{w}_{\text{opt}}^H \mathbf{\Sigma}_z \mathbf{w}_{\text{opt}}},$$

where $\mathbf{\Sigma}_z = \mathbb{E}\big(\mathbf{z}(t)\, \mathbf{z}^H(t)\big)$ is the covariance of the interference plus noise.

Since the purpose of the Nyström beamformer is to achieve satisfactory error performance for a low computational cost, we also compared the running time of each algorithm. For the low-SNR case (-10 dB), Figure 3 shows the average CPU

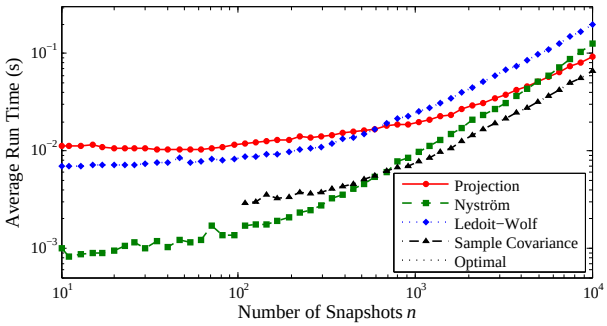

Figure 3. Average running time in seconds of various beamformers for INR = 20 dB and SNR = −10 dB, as a function of the number of snapshots.

time required to compute each beamformer on a personal computer with an Intel dual-core 2.33-GHz processor, as a function of the number of snapshots.

In the low-SNR case, the SINR performance of the Nyström beamformer is comparable to that of the projection beamformer, with the former trailing the latter by a margin of 0.4 to 1.6 dB. Both low-rank methods perform considerably better than beamforming using the sample and Ledoit-Wolf covariance estimators. We see that the sample covariance beamformer (which is undefined until the sample covariance becomes invertible at $n = 100$ snapshots) exhibits the poorest performance, due to the over-dispersion of its eigenvalues. Although shrinkage provided by the Ledoit-Wolf beamformer does improve the conditioning of the sample covariance beamformer, both approaches remain inferior to the low-rank methods until the number of snapshots grows large ($n > 4000$).

We observe similar performance trends in the medium-SNR (10 dB) and high-SNR (30 dB) cases. At an SNR of 10 dB, the SINR performance of the Nyström beamformer lags behind that of the projection beamformer by a margin of 0.06 to 1.4 dB, and at an SNR of 30 dB, the difference is minor (less than 0.15 dB) across all values of n. In both cases, the SINR performance of the Ledoit-Wolf and sample covariance estimators is consistently about 10 to 20 dB less than that of the low-rank methods.

In terms of computation, the Nyström beamformer requires about an order of magnitude less time to run than the projection and Ledoit-Wolf methods, for n up to around 100. However, as n grows large with p fixed, the complexity of all four approaches is dominated by the $O(n)$ cost of computing covariance terms over the set of snapshots, and thus the computational differences between the beamformers become less dramatic. Alternatively, if we were to let p grow while keeping n fixed, the computational cost would instead be dominated by the low-rank approximation step (or by matrix inversion, in the case of the Ledoit-Wolf or sample covariance beamformers). In this case, the Nyström beamformer would continue to exhibit significant computational savings when compared to the other methods.

VI. Example Application: Image Denoising

For our second example of an application for Nyström covariance estimation, we consider the problem of *image denoising* [29]–[31]. Let $\mathbf{z} \in \mathbb{R}^m$ represent an 8-bit grayscale image with m pixels, where each element $z_i \in \{0, \ldots, 255\}$ is the intensity of the i-th image pixel for $i = 1, \ldots, m$. Assume that we obtain a noisy version of this image $\mathbf{x} = \mathbf{z} + \mathbf{n}$, where $\mathbf{n} \sim \mathcal{N}_m\left(\mathbf{0}, \sigma^2 \mathbf{I}_m\right)$. Given $\mathbf{x}$, we wish to compute an estimate $\hat{\mathbf{z}}$ of the clean image $\mathbf{z}$.

Many approaches to image denoising involve computing local decompositions of the noisy signal over sets of nearby pixels (or "image patches"). We will investigate a denoising solution based on principal components analysis (PCA) of groups of patches, which requires estimating and then decomposing local covariance matrices.

Subspace Estimation for Image Denoising

We begin by developing a simple image model allowing us to perform denoising as a subspace estimation problem. Since natural images often possess a high degree of local similarity, a common assumption in image processing is that given a spatially proximate set of pixels, there exists some transformation under which these pixels admit a sparse representation [29], [32]. Let us partition $\mathbf{z}$ into q sub-vectors $\mathbf{z}_1, \ldots, \mathbf{z}_q$, where each $\mathbf{z}_i \in \mathbb{R}^p$ is a local set of p pixels (referred to as a "patch"), with $m = p\,q$. The observed image patches are $\mathbf{x}_i = \mathbf{z}_i + \mathbf{n}_i$, where $\mathbf{n}_i \sim \mathcal{N}_p\left(\mathbf{0}, \sigma^2 \mathbf{I}_p\right)$ for $i = 1, \ldots, q$. Assume now that each patch $\mathbf{z}_i$ is restricted to a subspace of dimension $k_i \leq p$, which we denote S_i. In this case, it can be shown [33] that a linear least-squares estimator of $\mathbf{z}_i$ given $\mathbf{x}_i$ is

$$\hat{\mathbf{z}}_i \equiv \mathbf{P}_i \mathbf{x}_i = \mathbf{z}_i + \mathbf{P}_i \mathbf{n}_i \tag{18}$$

where $\mathbf{P}_i$ represents an orthogonal projection onto S_i. This estimator preserves $\mathbf{z}_i$, while removing all noise except for the component in the signal subspace.

In practical applications, the subspaces $S_1, \ldots, S_q$ typically are not known and must be estimated from the noisy image. Consider a set of n noisy patches $\{\mathbf{x}_i : i \in I\}$, which lie within a region of the image defined by a set of n patch indices $I \subseteq \{1, \ldots, q\}$. If the patches in this region have similar signal characteristics, then we may assume that all S_i are equal for $i \in I$. This assumption allows us to estimate the subspaces from the k_i principal components of the sample covariance matrix

$$\mathbf{S} = \frac{1}{n} \sum_{i \in I} \mathbf{x}_i \mathbf{x}_i^T.$$

Note that in practice the subspace dimension is also unknown and may vary across regions. To address this problem, one may attempt to solve the rank estimation problem of determining k_i from the noisy image. However, for our simple example we set k_i to a fixed value across all images.

By repeating the component analysis for all regions, we can obtain a full set of projection matrices for computing the image estimate. We will refer to image estimation as in (18) where the projections are estimated from the principal components of sample covariances as the *PCA image denoiser*.

Depending on the size of the image, estimating the full set of orthogonal projections for PCA denoising can require computing and then decomposing or inverting a large number

Figure 4. High-resolution test images for image denoising experiments. Top left: *flower* (2256×1504). Top right: *leaves* (3008×2000). Bottom left: *hdr* (3072×2048). Bottom right: *deer* (4032×2640).

of covariance matrices. Consequently, we may realize significant improvements in computation by replacing the sample covariance with a Nyström covariance estimate, which we will refer to as the *Nyström image denoiser*.

Although we have defined both approaches for disjoint image patches and regions, one may want to allow for some amount of overlap among these sets. This modification increases the number of available samples, while also mitigating some of the artifacts that occur at patch boundaries. While commonly used in practice [29], [30], note that allowing for overlapping patches conflicts with independence assumptions of our additive noise model. It also increases computation due to the additional number of patches, and because estimated patches may need to be reweighted before they are combined.

Experimental Results

To compare the PCA and Nyström image denoisers, we examined their performance when applied to a selection of 8-bit high-resolution test images from [34]. The four images used are shown in Figure 4. Our primary measure of performance is the *peak signal-to-noise ratio* (PSNR) of the estimated image, a standard metric used throughout the image processing literature. The PSNR of the denoised image $\hat{\mathbf{z}}$ is defined as

$$\text{PSNR} = \frac{z_{\max}^2}{\|\hat{\mathbf{z}} - \mathbf{z}\|^2},$$

where $z_{\max}$ denotes the maximum allowable pixel value (in our case, 255). After generating noisy versions of each image for noise levels of $\sigma = 10$, 20, and 50, we attempted to reconstruct the original image using the PCA and Nyström approaches, assuming a fixed subspace dimension of $k = 4$.

By examining denoising results for different patch and region sizes over the selection of test images, we determined a set of default parameter values that yielded a reasonable balance of performance and computation. For both algorithms, we divided the image into regions of 32×32 pixels, with adjacent regions having 50% overlap. We then divided each region into 8×8 patches, also with 50% overlap. Thus, each region contained $n = 49$ patches, which were used to estimate a local covariance of dimension $p = 64$.

In the case of the Nyström denoiser, covariance estimation was performed for each region using the Nyström covariance estimator, conditioned on a set of k patch vectors chosen uniformly at random. Once the k principal components and the estimated projection were computed, the denoised patches were superimposed to reconstruct an estimate of the original image.

As a benchmark, we provide results from two other patch-based denoising methods: the K-SVD algorithm of [30] and the BM3D algorithm of [31]. The first algorithm performs denoising based on a trained "dictionary" of components obtained from the noisy image, while the second jointly filters groups of image patches by arranging them into 3-D arrays. For both algorithms, denoising was performed using code from the authors' respective websites, with most parameters set to their default values. The only parameters we adjusted were the dictionary size and maximum number of training blocks for the K-SVD algorithm; to accommodate the high resolution of

Table I
AVERAGE PSNR (IN DB) OF DENOISING ALGORITHMS FOR HIGH-RESOLUTION TEST IMAGES

Image	σ / PSNR	PCA	Nyström	K-SVD	BM3D
flower	10 / 28.13	37.81	38.80	34.84	43.76
	20 / 22.11	32.31	34.02	32.18	40.47
	50 / 14.15	24.50	26.69	28.73	36.06
leaves	10 / 28.13	27.99	27.26	31.39	35.39
	20 / 22.11	27.00	26.15	27.56	31.61
	50 / 14.15	22.92	22.88	23.16	26.78
hdr	10 / 28.13	37.18	37.54	32.58	42.51
	20 / 22.11	32.08	33.42	29.13	39.20
	50 / 14.15	24.44	26.55	25.42	34.76
deer	10 / 28.13	33.47	33.47	30.51	34.19
	20 / 22.11	30.58	31.46	26.67	33.33
	50 / 14.15	24.17	26.05	22.57	31.98

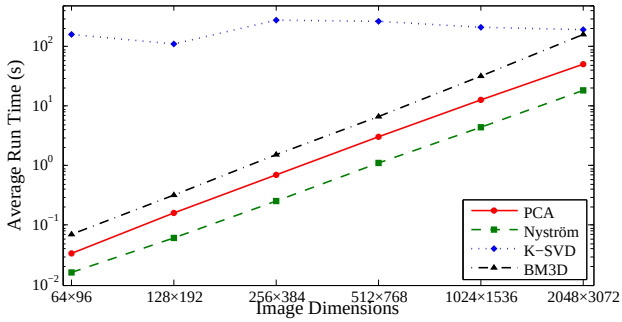

Figure 5. Average running times in seconds of various denoising algorithms when applied to different sizes of *hdr* image.

test images, these values were increased to 1024 and 130,000, respectively.

Results for all four algorithms are listed in Table I. For each test image, algorithm, and noise level, we list the average empirical PSNR over 10 realizations of each noisy image. We see that the performance of the PCA and Nyström denoisers is comparable to the benchmark algorithms, with average PSNR values typically falling somewhere between those of the K-SVD and BM3D algorithms. These results suggest that subspace-based denoising provides a reasonable venue for testing the capabilities of the Nyström covariance estimator.

In comparing the PCA and Nyström denoisers, we find that despite requiring significantly less computation, the Nyström approach actually performs slightly better in most cases. One explanation for this behavior is that the shrinkage performed by the Nyström covariance estimator allows for improved subspace estimation.

To better illustrate the computational differences between the algorithms, in Figure 5 we show the average run time of each method when applied to a noisy version of the *hdr* image with $\sigma = 20$. Results are shown for the original 2048×3072 image, as well as for resampled versions with sizes of 64×96, 128×192, 256×384, 512×768, and 1024×1536. We see that the costs of the PCA, Nyström and BM3D denoisers scale similarly with image size, with the Nyström denoiser performing about 2–3 times faster than the PCA denoiser and 3–9 times faster than BM3D. Note that due to the complexity of its training step, the cost of K-SVD is high (though relatively constant) across all image sizes.

In Figure 6, we show 256×256 close-ups of denoising results for the test image *hdr*, given a noise level of $\sigma = 20$. The Nyström-denoised image is visibly smoother and contains fewer artifacts than does the PCA-denoised image, as there is less mismatch between its estimated subspace and the "true" subspace represented by the clean image data. This difference is reflected in the PSNR values for each method.

VII. SUMMARY

In this article, we developed the Nyström approximation as a low-rank covariance estimator. In addition to deriving expressions for its bias and mean squared error for the case of normally-distributed data, we showed that the Nyström covariance estimator shrinks the sample eigenvalues. This shrinkage allows the estimator to achieve performance that is comparable to (or at times, better than) the sample covariance, particularly in cases where the number of samples is less than the dimension of the data. Moreover, because of the computational advantages of the Nyström covariance estimator, its eigenvalues and eigenvectors can be computed for far less than the cost of spectral analysis of the sample covariance.

We illustrated the potential of the Nyström covariance estimator through its use in two example applications: array signal processing and image denoising. In the first example, we adapted a projection-based beamforming algorithm to utilize the Nyström estimator, resulting in reduced computation with little degradation in our ability to recover the desired signal. In the second example, we developed a simple PCA-based algorithm for image denoising, and then showed how the Nyström covariance estimator could be used to reduce computation while maintaining or even improving denoising performance.

REFERENCES

[1] J. F. Cardoso and A. Souloumiac, "Blind beamforming for non-Gaussian signals," *IEE Proceedings F: Radar and Signal Processing*, vol. 140, no. 6, pp. 362–370, 1993.

[2] Y. Ephraim and H. L. Van Trees, "A signal subspace approach for speech enhancement," *IEEE Transactions on Speech and Audio Processing*, vol. 3, no. 4, pp. 251–266, 1995.

[3] A. Belouchrani, A.-M. Karim, J. F. Cardoso, and E. Moulines, "A blind source separation technique using second-order statistics," *IEEE Transactions on Signal Processing*, vol. 45, no. 2, pp. 434–444, 1997.

[4] C. Lanczos, "An iteration method for the solution of the eigenvalue problem of linear differential and integral operators," *Journal of Research of the National Bureau of Standards*, vol. 45, no. 4, pp. 255–282, 1950.

[5] W. E. Arnoldi, "The principle of minimized iterations in the solution of the matrix eigenvalue problem," *Quarterly of Applied Mathematics*, vol. 9, pp. 17–29, 1951.

[6] V. A. Marčenko and L. A. Pastur, "Distribution of eigenvalues for some sets of random matrices," *Mathematics of the USSR: Sbornik*, vol. 1, pp. 457–483, 1967.

[7] I. M. Johnstone, "On the distribution of the largest eigenvalue in principal components analysis," *The Annals of Statistics*, vol. 29, no. 2, pp. 295–327, 2001.

[8] C. Stein, "Estimation of a covariance matrix," Reitz Lecture, 39th Annual Meeting of the Institute of Mathematical Statistics, Atlanta, GA, 1975.

[9] L. R. Haff, "Empirical Bayes estimation of the multivariate normal covariance matrix," *The Annals of Statistics*, vol. 8, no. 3, pp. 586–597, 1980.

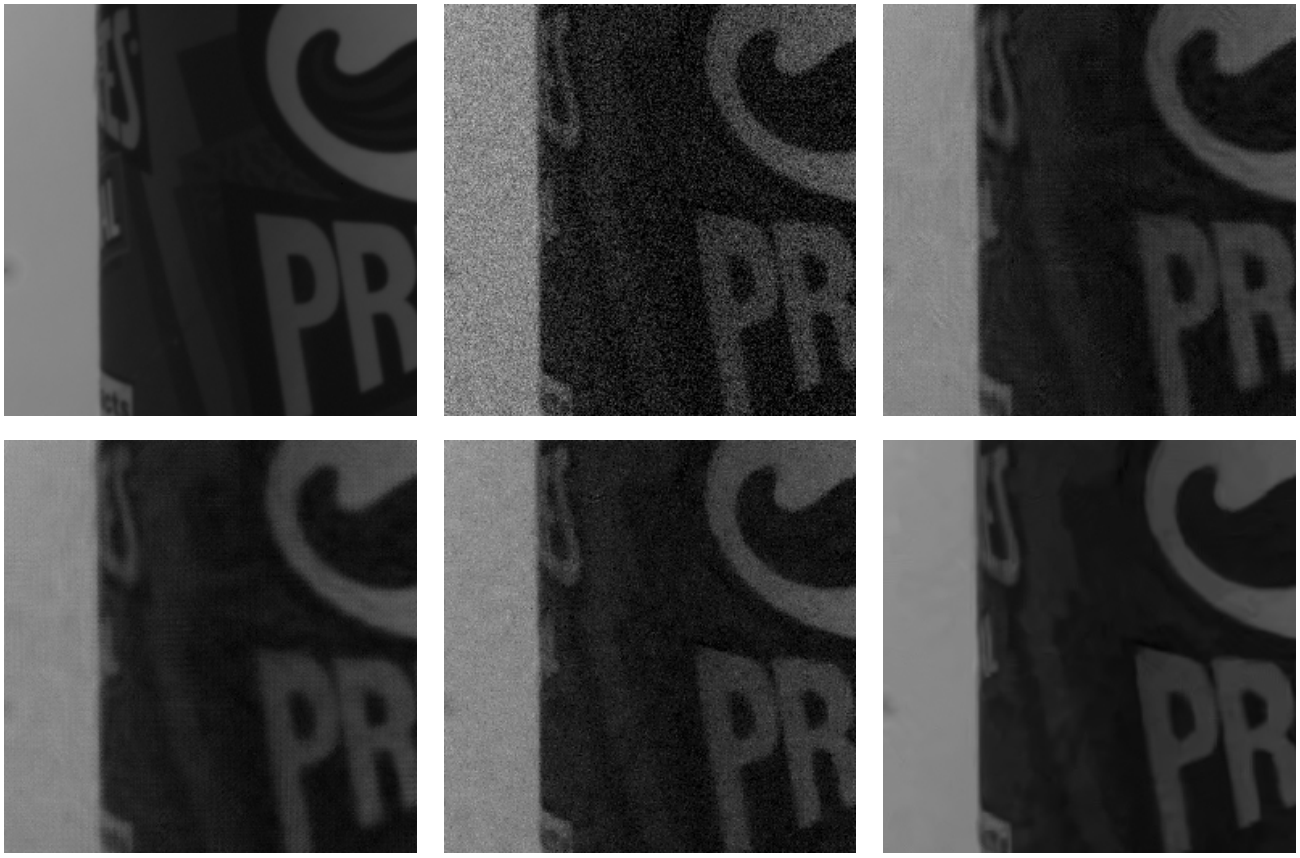

Figure 6. Close-up of denoising results for test image *hdr*, given $\sigma = 20$. Top left: original image. Top right: noisy image (PSNR = 22.11 dB, SSIM = 0.947). Middle left: PCA (PSNR = 32.79 dB, SSIM = 0.975). Middle right: Nyström (PSNR = 34.35, SSIM = 0.976). Bottom left: K-SVD (PSNR = 29.12 dB, SSIM = 0.970). Bottom right: BM3D (PSNR = 39.22 dB, SSIM = 0.986).

[10] D. K. Dey and C. Srinivasan, "Estimation of a covariance matrix under Stein's loss," *The Annals of Statistics*, vol. 13, no. 4, pp. 1581–1591, 1985.

[11] O. Ledoit and M. Wolf, "A well-conditioned estimator for large-dimensional covariance matrices," *Journal of Multivariate Analysis*, vol. 88, no. 2, pp. 365–411, 2004.

[12] C. K. I. Williams and M. Seeger, "Using the Nyström method to speed up kernel machines," in *Advances in Neural Information Processing Systems*, 2000.

[13] C. Fowlkes, S. Belongie, F. Chung, and J. Malik, "Spectral grouping using the Nyström method," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 2, pp. 1–12, 2004.

[14] N. Arcolano and P. J. Wolfe, "Estimating principal components of large covariance matrices using the Nyström method," in *Proceedings of the IEEE Conference on Acoustics, Speech, and Signal Processing*, 2011.

[15] P. L. Leung and R. J. Muirhead, "Estimation of parameter matrices and eigenvalues in manova and canonical correlation," *The Annals of Statistics*, vol. 15, no. 4, pp. 1651–1666, 1987.

[16] T. W. Anderson, "Asymptotic theory for principal component analysis," *The Annals of Mathematical Statistics*, vol. 34, no. 1, pp. 122–148, 1963.

[17] P. Drineas and M. W. Mahoney, "On the Nyström method for approximating a Gram matrix for improved kernel-based learning," *The Journal of Machine Learning Research*, vol. 6, pp. 2153–2175, 2005.

[18] M.-A. Belabbas and P. J. Wolfe, "Spectral methods in machine learning: new strategies for very large data sets," in *Proceedings of the National Academy of the Sciences of the United States of America*, vol. 106, 2009, pp. 369–374.

[19] A. K. Gupta and D. K. Nagar, *Matrix Variate Distributions*. Boca Raton: Chapman & Hall / CRC, 2000.

[20] R. Bhatia, *Matrix Analysis*. New York: Springer, 1997, vol. 169.

[21] L. E. Brennan and I. S. Reed, "Theory of adaptive radar," *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-9, no. 2, pp. 237–252, 1973.

[22] B. D. Van Veen and K. M. Buckley, "Beamforming: a versatile approach to spatial filtering," *IEEE ASSP Magazine*, vol. 5, no. 2, pp. 4–24, 1988.

[23] D. D. Feldman and L. J. Griffiths, "A projection approach to robust adaptive beamforming," *IEEE Transactions on Signal Processing*, vol. 42, no. 4, pp. 867–876, 1994.

[24] K. L. Bell, Y. Ephraim, and H. L. Van Trees, "A Bayesian approach to robust adaptive beamforming," *IEEE Transactions on Signal Processing*, vol. 48, no. 2, pp. 386–398, 2000.

[25] H. L. Van Trees, *Detection, Estimation, and Modulation Theory, Part I*. New York: John Wiley & Sons, 1968.

[26] P. Parker, P. J. Wolfe, and V. Tarokh, "A signal processing application of randomized low-rank approximations," in *IEEE Workshop on Statistical Signal Processing*, 2005.

[27] D. N. Spendley and P. J. Wolfe, "Adaptive beamforming using fast low-rank covariance matrix approximations," in *Proceedings of the IEEE Radar Conference*, 2008, pp. 1–5.

[28] P. A. Zulch, J. S. Goldstein, J. R. Guerci, and I. S. Reed, "Comparison of reduced-rank signal processing techniques," in *Proceedings of the Asilomar Conference on Signals, Systems & Computers*, 1998, pp. 421–245.

[29] D. D. Muresan and T. W. Parks, "Adaptive principal components and image denoising," in *Proceedings of the International Conference on Image Processing*, 2003, pp. 101–104.

[30] M. Elad and M. Aharon, "Image denoising via sparse and redundant representations over learned dictionaries," *IEEE Transactions on Image Processing*, vol. 15, no. 12, pp. 3736–3745, 2006.

[31] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3-D transform-domain collaborative filtering," *IEEE Transactions on Image Processing*, vol. 16, no. 8, pp. 2080–2095, 2007.

[32] S. Mika, B. Schölkopf, A. J. Smola, K.-R. Müller, M. Scholz, and G. Rätsch, "Kernel PCA and de-noising in feature spaces," in *Advances in Neural Information Processing Systems*, 1999.

[33] L. L. Scharf and C. Demeure, *Statistical Signal Processing: Detection, Estimation, and Time Series Analysis*. Reading, Massachusetts: Addison-Wesley, 1991.

[34] S. Garg, "The new test images," 2010. [Online]. Available: http://www.imagecompression.info/test_images/